

Cauchy-Davenport Theorem for linear maps: Simplification and Extension

John Kim* Aditya Potukuchi†

December 30, 2016

Abstract

We give a new proof of the Cauchy-Davenport Theorem for linear maps given by Herdade et al., (2015) in [2]. This theorem gives a lower bound on the size of the image of a linear map on a grid. Our proof is purely combinatorial and offers a partial insight into the range of parameters not handled in [2].

1 Introduction

Let \mathbb{F}_p be the field containing p elements, where p is a prime, and let $A, B \subseteq \mathbb{F}_p$. The Cauchy-Davenport Theorem gives a lower bound on the size of the sumset $A + B \stackrel{\text{def}}{=} \{a + b \mid a \in A, b \in B\}$ (for more on sumsets, see, for example, [3]). The size of the sumset can be thought of as the size of the image of the linear map $(x, y) \rightarrow x + y$, where $x \in A$, and $y \in B$. Thus the theorem can be restated as follows:

Theorem 1.1 (Cauchy-Davenport Theorem). *Let p be a prime, and $L : \mathbb{F}_p \times \mathbb{F}_p \rightarrow \mathbb{F}_p$ be a linear map that takes (a, b) to $a + b$. For $A, B \subseteq \mathbb{F}_p$, Let $L(A, B)$ be the image of L on $A \times B$. Then,*

$$|L(A, B)| \geq \min(|A| + |B| - 1, p)$$

In [2], this notion was extended to study the sizes of images of general linear maps on product sets. A lower bound was proved using the polynomial method (via a nonstandard application of the Combinatorial Nullstellensatz [1]). In this paper, we give a simpler, and combinatorial proof of the same using just the Cauchy-Davenport Theorem.

*Department of Mathematics, Rutgers University. Research supported in part by NSF Grant Number DGE-1433187. jonykim@math.rutgers.edu.

†Department of Computer Science, Rutgers University. aditya.potukuchi@cs.rutgers.edu.

Notation: For a linear map $L : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^m$, and for $S_1, S_2, \dots, S_n \subseteq \mathbb{F}_p$, we use $L(S_1, S_2, \dots, S_n)$ to denote the image of L on $S_1 \times S_2 \times \dots \times S_n$. The *support* of a vector is the set of nonzero entries in the vector. A *min-support vector* in a set V of vectors is a nonzero vector of minimum support size in V .

Theorem 1.2 (Main Theorem). *Let p be a prime, and $L : \mathbb{F}_p^{m+1} \rightarrow \mathbb{F}_p^m$ be a linear map of rank m . Let $A_1, A_2, \dots, A_{m+1} \subseteq \mathbb{F}_p$ with $|A_i| = k_i$. Further, suppose that $\min_i(k_i) + \max_i(k_i) < p$. Let S be the support of $\ker(L)$, and $S' = [m+1] \setminus S$. Then*

$$|L(A_1, A_2, \dots, A_{m+1})| \geq \left(\prod_{j \in S'} k_j \right) \cdot \left(\prod_{i \in S} k_i - \prod_{i \in S} (k_i - 1) \right)$$

As noted in [2], this bound is tight for every m and p . We restrict our theorem to study only maps from \mathbb{F}_p^{m+1} to \mathbb{F}_p^m of rank m for two reasons mainly: (1) It is simpler to state, and contains the tight case and (2) We are unable to prove any better bounds if the rank is not m . It is not clear to us what the correct bound for the general case is.

We also show the following result for the size of the image for certain full rank linear maps from $\mathbb{F}_p^n \rightarrow \mathbb{F}_p^{n-1}$ when the size of the sets it is evaluated on are all large enough.

Theorem 1.3. *Let $L : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^{n-1}$ be a linear map given by $L(x_1, \dots, x_n) = (x_1 + x_n, x_2 + x_n, \dots, x_{n-1} + x_n)$. Let $S_1, \dots, S_n \subseteq \mathbb{F}_p$ with $|S_i| = k$ for $i \in [n]$ such that $k > \frac{(n-1)p}{n}$, then $|L(S_1, \dots, S_n)| = p^{n-1}$ (i.e., $L(S_1, \dots, S_n) = \mathbb{F}_p^{n-1}$).*

The theorems do not, however, give tight bounds for all set sizes, for example if $\min_i |A_i| > p/2$. It would be interesting to obtain a tight bound even for the simple linear map $(x, y, z) \rightarrow (x+z, y+z)$ on the product set $A_1 \times A_2 \times A_3 \subseteq \mathbb{F}_p^3$ which holds for all sizes of the A_i 's.

2 The Theorem

2.1 The Main Lemma

The idea is that since the size of the image is invariant under row operations of L , we perform row operations to isolate a ‘hard’ part, which gives the main part of the required lower bound

Our proof proceeds by induction on the dimension of the linear map. The base case is given by the Cauchy Davenport Theorem.

Lemma 2.1. *Let $L : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^{n-1}$ be a linear map such that $L(x_1, \dots, x_n) = (x_1 + x_n, x_2 + x_n, \dots, x_{n-1} + x_n)$. Let $S_1, \dots, S_n \subseteq \mathbb{F}_p$ with $|S_i| = s_i$ such that $\min_i(s_i) + \max_i(s_i) \leq p + 1$. Then $|L(S_1, \dots, S_n)| \geq \prod_{i=1}^n s_i - \prod_{i=1}^n (s_i - 1)$*

Proof. We use the shorthand notation $|L| \stackrel{\text{def}}{=} |L(S_1, S_2 \dots S_n)|$. W.L.O.G, let S_1 be such that $|S_1| = \min_{i \in [n-1]} (|S_i|)$.

A preliminary observation is that $|S_1| + |S_n| \leq p + 1$, and therefore, by the Cauchy-Davenport Theorem,

$$|S_1 + S_n| \geq s_1 + s_n - 1 \quad (1)$$

The proof proceeds by induction on n . If $n = 2$, the result $|L| \geq s_1 \cdot s_2 - (s_1 - 1) \cdot (s_2 - 1) = s_1 + s_2 - 1$ is given by the Cauchy-Davenport Theorem.

For every $a \in \mathbb{F}_p$, we have $T_a \stackrel{\text{def}}{=} \{x_n \in S_n \mid \exists x_1 \in S_1, x_1 + x_n = a\}$, and $t_a \stackrel{\text{def}}{=} |T_a|$. We now look at the restricted linear map $L|_{x_1+x_n=a}$. In this case, the induction is on sets $S_2, \dots, S_{n-1} \times T_a$. This is equivalent to restricting S_n to the set T_a , and dropping S_1 , since for every $x_n \in S_n$, there is a unique $x_1 \in S_1$ such that $x_1 + x_n = a$.

We first observe that the conditions are satisfied, i.e., $\min_i (|S_i|) + \max_i (|S_i|) \leq p + 1$, since $t_a \leq \min(|S_1|, |S_n|)$. Also the resulting linear map is of the same form, i.e., $L|_{x_1+x_n=a}(x_2, \dots, x_n) = (x_2 + x_n \dots x_{n-1} + x_n)$. (In reality, $L|_{x_1+x_n=a}$ is a map from \mathbb{F}_p^n to \mathbb{F}_p^{n-1} , given by $L|_{x_1+x_n=a}(x_1, x_2, \dots, x_n) = (a, x_2 + x_n \dots x_{n-1} + x_n)$ but we drop the first coordinate because it is fixed, i.e., a)

By induction hypothesis, the number of points in the image of $L|_{x_1+x_n=a}$ is at least:

$$\left(\prod_{i=2}^{n-1} s_i \right) t_a - \left(\prod_{i=2}^{n-1} (s_i - 1) \right) (t_a - 1)$$

Summing over all $a \in \mathbb{F}_p$, we get a bound on the number of points in the image:

$$\begin{aligned} |L| &\geq \sum_{a \in \mathbb{F}_p, t_a \neq 0} \left(\left(\prod_{i=2}^{n-1} s_i \right) t_a - \left(\prod_{i=2}^{n-1} (s_i - 1) \right) (t_a - 1) \right) \\ &= \left(\prod_{i=2}^{n-1} s_i \right) \sum_{a \in \mathbb{F}_p} t_a - \left(\prod_{i=2}^{n-1} (s_i - 1) \right) \sum_{a \in \mathbb{F}_p, t_a \neq 0} (t_a - 1) \\ &\geq \prod_{i=1}^n s_i - \prod_{i=1}^n (s_i - 1) \end{aligned}$$

The last inequality comes from observing that $\sum_{a \in \mathbb{F}_p} t_a = s_1 s_n$, and an upper bound on $\sum_{a \in \mathbb{F}_p, t_a \neq 0} (t_a - 1)$, by using 1. We have $\sum_{a \in \mathbb{F}_p, t_a \neq 0} (t_a - 1) = \sum_{a \in \mathbb{F}} t_a - \sum_{a \in \mathbb{F}_p} \mathbb{1}_{t_a \neq 0} = \sum_{a \in \mathbb{F}} t_a + |S_1 + S_n| \leq s_1 s_n - (s_1 + s_n - 1)$. \square

2.2 Arriving at the Main Theorem

The first step in arriving at the main theorem is exactly as in [2]. For completeness, we describe it here. The idea is to transform a general linear map into a specific form, without reducing the size of the image (in fact, here it remains the same). This step is very intuitive, but describing it requires some setup.

Let $L : \mathbb{F}_p^{m+1} \rightarrow \mathbb{F}_p^m$ be an \mathbb{F}_p -linear map of rank m . Let v be a non-zero min-support vector of $\ker(L)$. So, we have $Lv = 0$. The main observation is that under row operations, two quantities remain unchanged: the size of the image of L , and the size of the support of the min-support vector in the kernel.

Let r_1, \dots, r_m be the rows, and c_1, c_2, \dots, c_{m+1} be the columns of associated to L with respect to the standard basis. We show that one can perform elementary row operations, and some column operations on L while preserving the size of the image.

Lemma 2.2. *The size of the image of L does not change under*

1. *Elementary row operations.*
2. *Scaling any column c_i by some $d \in \mathbb{F}_p \setminus \{0\}$ and scaling every element of A_i by d .*
3. *Swapping any two columns c_i and c_j , and swapping sets A_i and A_j .*

Proof. We prove this by considering each given operation separately.

1. Suppose L' was obtained from L by elementary row operations. There is an invertible linear map M such that $M \cdot L = L'$. This gives the bijection from every vector v in the image of L , to the vector $M \cdot v$ in the image of L' .
2. Suppose L' was obtained from L by scaling column c_i by $d \in \mathbb{F}_p \setminus \{0\}$, and scaling the set A_i by d^{-1} . We map every vector $(u_1, \dots, u_m) \in L(A_1, \dots, A_i, \dots, A_{m+1})$, to the vector $(u_1, \dots, u_m) \in L'(A_1, \dots, d^{-1} \cdot A_i, \dots, A_{m+1})$. Here $d^{-1} \cdot A_i \stackrel{\text{def}}{=} \{d^{-1}a_i \mid a_i \in A_i\}$. This map is invertible.
3. Suppose L' was obtained from L by switching columns c_i and c_j , and swapping the sets A_i and A_j . We map every vector $(u_1, \dots, u_m) \in L(A_1, \dots, A_i, \dots, A_j, \dots, A_{m+1})$ to the identical vector $(u_1, \dots, u_m) \in L'(A_1, \dots, A_j, \dots, A_i, \dots, A_{m+1})$. This map is invertible.

For every given operation, we have a bijection between the images of L before and after the operation.

□

Observation 2.3. *After the operations stated in Lemma 2.2, the size of the support of the min-support vector in $\ker(L)$ does not change.*

To see this, we first observe that the kernel has rank 1, and is orthogonal to the row span of L . Therefore, all nonzero vectors in $\ker(L)$ have the same support. Since, row operations do not change the row span of L , the resulting kernel spans the same subspace of \mathbb{F}^{m+1} , and therefore, the size of the support of the vectors in $\ker(L)$ does not change.

Next, we do the following operations, each of which preserves the size of the image.

1. Perform row operations so that the last m columns form an identity matrix.
2. Scale the rows so that the first column of every row is 1.
3. Scale the last m columns so that every nonzero entry in L is 1.

After we perform these operations, we have a linear map where the first column consists of 1's and 0's and the remaining m columns form an identity matrix. Let the S' be the set of indices of rows containing 1's in the first column. Consider the vector $v = -e_1 + \sum_{i \in S'} e_{i+1}$. This vector has support $|S'| + 1$, and lies in the kernel of L . Therefore, $|S| = |S'| + 1$.

Proof of Theorem 1.2. Apply the transformation from Lemma 2.2 to L to reduce it to the simple form. Let S' be the set of rows where the first column is nonzero. Consider the restriction of L on the the coordinates given by S . By Lemma 2.1, the size of this image is at least $(\prod_{i \in S} k_i - \prod_{i \in S} (k_i - 1))$.

The linear map restricted to the coordinates $[m] \setminus S$ is nothing but the identity map, so the size of the image is $\prod_{i \notin S} |A_i|$, and is independent of the linear map restricted to S . Putting them together, we have the desired result. \square

3 The case when $2k > p + 1$

The proof of Lemma 2.1 breaks down when $s_1 + s_n > p + 1$ and, unfortunately, we do not know how to fix this issue. Consider, for example, the simplest nontrivial case where $m = 2$, i.e., $L(x, y, z) = (x + z, y + z)$, and we are interested in the size of the image of L on $X \times Y \times Z$, further suppose, for simplicity, that $|X| = |Y| = |Z| = k$. If $k < \frac{p+1}{2}$, then the above bound holds, and is tight. If $k > \frac{2p}{3}$, then L covers \mathbb{F}_p^2 , i.e., $|L(A, B, C)| = p^2$. This makes the case in between the interesting one. We conjecture that the correct lower bound is the size of the image of L when $X = Y = Z = \{1, 2, \dots, k\}$. Towards this, we are able to prove a partial result (Lemma 3.2) using the above method.

We will need the following Lemma:

Lemma 3.1. *Let $X, Y \subseteq \mathbb{F}_p$ and $t_a = |\{(x, y) \in X \times Y : x + y = a\}|$. Then for every $a \in \mathbb{F}_p$:*

$$|X| + |Y| - p \leq t_a \leq \min(|X|, |Y|).$$

Proof. The bounds follow from the fact that t_a can be written as the size of the intersection of two sets of sizes $|X|$ and $|Y|$:

$$t_a = |X \cap (a - Y)|.$$

□

Now we state the partial result:

Theorem 3.2. *Let $L : \mathbb{F}_p^3 \rightarrow \mathbb{F}_p^2$ be the linear map defined by $L(x, y, z) = (x + z, y + z)$. Let $X, Y, Z \subset \mathbb{F}_p$ be sets of size k , where $k \geq \frac{p+1}{2}$. Then we have the following lower bound:*

$$|L(X, Y, Z)| \geq \min(p^2 + 3k^2 - (2p + 1)k, p^2).$$

Proof. Let $T_a \stackrel{\text{def}}{=} \{z \in Z \mid \exists x \in X, x + z = a\}$, with $t_a \stackrel{\text{def}}{=} |T_a|$. Looking at this restriction, $L|_{x+z=a}$, by Cauchy-Davenport Theorem, there are at least $\min(t_a + k - 1, p)$ points of $L(X, Y, Z)$ on $L_{x+z=a}(Y, T_a)$. By summing over all $a \in \mathbb{F}_p$, we get a lower bound on the size of $L(X, Y, Z)$:

$$\begin{aligned} |L(X, Y, Z)| &\geq \sum_{a \in \mathbb{F}_p} \min(t_a + k - 1, p) \\ &= \sum_{a \in \mathbb{F}_p} \min(t_a, p - k + 1) + p(k - 1) \\ &= \sum_{a: t_a \leq p - k + 1} t_a + \sum_{a: t_a > p - k + 1} (p - k + 1) + p(k - 1). \end{aligned}$$

We now want to remove the dependence of the lower bound on the t_a by considering the worst case scenario, where the t_a take values that minimize the lower bound. First, we observe $\sum_{a \in \mathbb{F}_p} t_a = k^2$, a fixed quantity. So to minimize the above lower bound for $|L(X, Y, Z)|$, we need t_a to be maximal for as many $a \in \mathbb{F}_p$ as possible.

By Lemma 3.1, we know that $2k - p \leq t_a \leq k$. We set $t_a = k$ for as many $a \in \mathbb{F}_p$ as possible, and the remainder of the $t_a = 2k - p$. This gives:

$$\begin{aligned} |L(X, Y, Z)| &\geq \sum_{a: t_a \leq p - k + 1} t_a + \sum_{a: t_a > p - k + 1} (p - k + 1) + p(k - 1) \\ &\geq k(2k - p) + (p - k)(p - k + 1) + p(k - 1) \\ &= 3k^2 + p^2 - (2p - 1)k. \end{aligned}$$

□

As a corollary, we get, independent of theorem 1.3, the following corollary:

Corollary 3.3. *If the linear map L , and the sets A, B, C were as above, with $|A| = |B| = |C| = k$, and $k > \frac{2p}{3}$, then $L(A, B, C) = p^2$.*

We would like to point out that at the two extremes, i.e., when $k = \frac{p+1}{2}$, and when $k = \lceil \frac{2p}{3} \rceil$, the above bound matches the ‘correct’ lower bound.

3.1 Proof of Theorem 1.3

We prove theorem 1.3 via a slightly stronger claim

Claim 3.4. *Let $L : \mathbb{F}_p^n \rightarrow \mathbb{F}_p^{n-1}$ be a linear map given by $L(x_1, \dots, x_n) = (x_1 + x_n, x_2 + x_n, \dots, x_{n-1} + x_n)$. Let $S_1, \dots, S_n \subseteq \mathbb{F}_p$ with $|S_i| = k$ for $i \in [n-1]$, and $|S_n| = k'$. Further, suppose that $(n-1)k + k' \geq (n-1)p + 1$, then $|L(S_1, \dots, S_n)| = p^{n-1}$.*

Proof. We prove this by induction on n , analogous to Lemma 2.1. The case where $n = 2$ is, again, given by the Cauchy-Davenport Theorem.

For $a \in \mathbb{F}_p$, $T_a \stackrel{\text{def}}{=} \{x_n \in S_n \mid \exists x_1 \in S_1, x_1 + x_n = a\}$ with $t_a \stackrel{\text{def}}{=} |T_a|$. Looking at this restriction of L (i.e., $x_1 + x_n = a$), we have a linear map, $L_{x_1+x_n=a}$ on the sets $S_2 \times S_3 \times \dots \times T_a$, given by the $L_{x_1+x_n=a}(x_2, \dots, x_n) = (x_2 + x_n, \dots, x_{n-1} + x_n)$. (similar to Lemma 2.1, we drop the first coordinate).

Here, $|S_i| = k$ for $i = 2, \dots, n-1$, and $|T_a| \geq k + k' - p$, by Lemma 3.1. Further, the required condition holds, i.e.,:

$$(n-2)k + t_a \geq (n-2)k + k + k' - p = (n-1)k + k' - p \geq (n-2)p + 1.$$

Therefore, by induction hypothesis $|L_{x_1+x_n=a}(S_2, \dots, S_{n-1}, T_a)| = p^{n-2}$. Since this holds for every $a \in \mathbb{F}_p$, we have $|L(S_1, \dots, S_n)| = p^{n-1}$. \square

In particular, Lemma 3.4 tells that for the linear map L given by $L(x_1, \dots, x_n) = (x_1 + x_n, x_2 + x_n, \dots, x_{n-1} + x_n)$ on $S_1 \times S_2 \times \dots \times S_n$, if $|S_i| \geq \frac{(n-1)p}{n}$, then $L(S_1, \dots, S_n) = \mathbb{F}_p^{n-1}$.

4 Acknowledgements

We would like to thank Swastik Kopparty for the discussions and the many helpful ideas.

References

- [1] N. Alon, *Combinatorial nullstellensatz*, Combinatorics, Probability and Computing, 8: 7-29, 1999.
- [2] S. Herdade, J. Kim, S. Kopparty, *A Cauchy-Davenport theorem for linear maps*, preprint, <http://arxiv.org/abs/1508.02100>.
- [3] T. Tao, V. Vu, *Additive combinatorics*, Cambridge University Press, 2006.